

ICALEPCS 2025

25 SEPTEMBER 2025
CHICAGO, IL, USA

Better software observability using Tango Controls with OpenTelemetry - experience at MAX IV

Anton Joubert (author)

Marcelo Alcocer (presenter)

Lin Zhu (co-author)

Benjamin Bertrand (co-author)



Agenda

Introduction

Examples

Performance and resources

Conclusion

Introduction

Observability

Understand complex system from outside
Logs, traces, metrics, profiles

Introduction

Observability

Understand complex system from outside
Logs, traces, metrics, profiles

OpenTelemetry

Observability framework
Vendor agnostic
Open source
Many programming languages

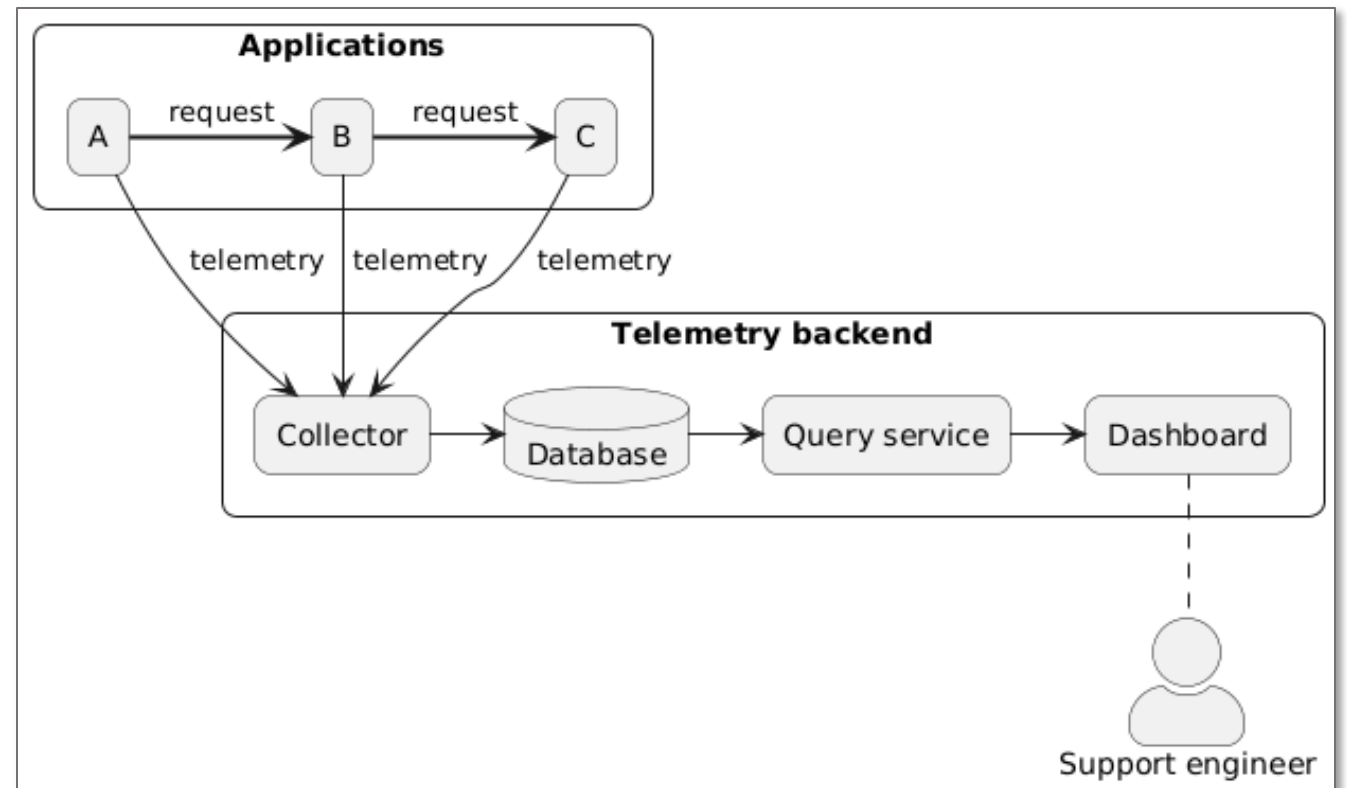
Introduction

Observability

Understand complex system from outside
Logs, traces, metrics, profiles

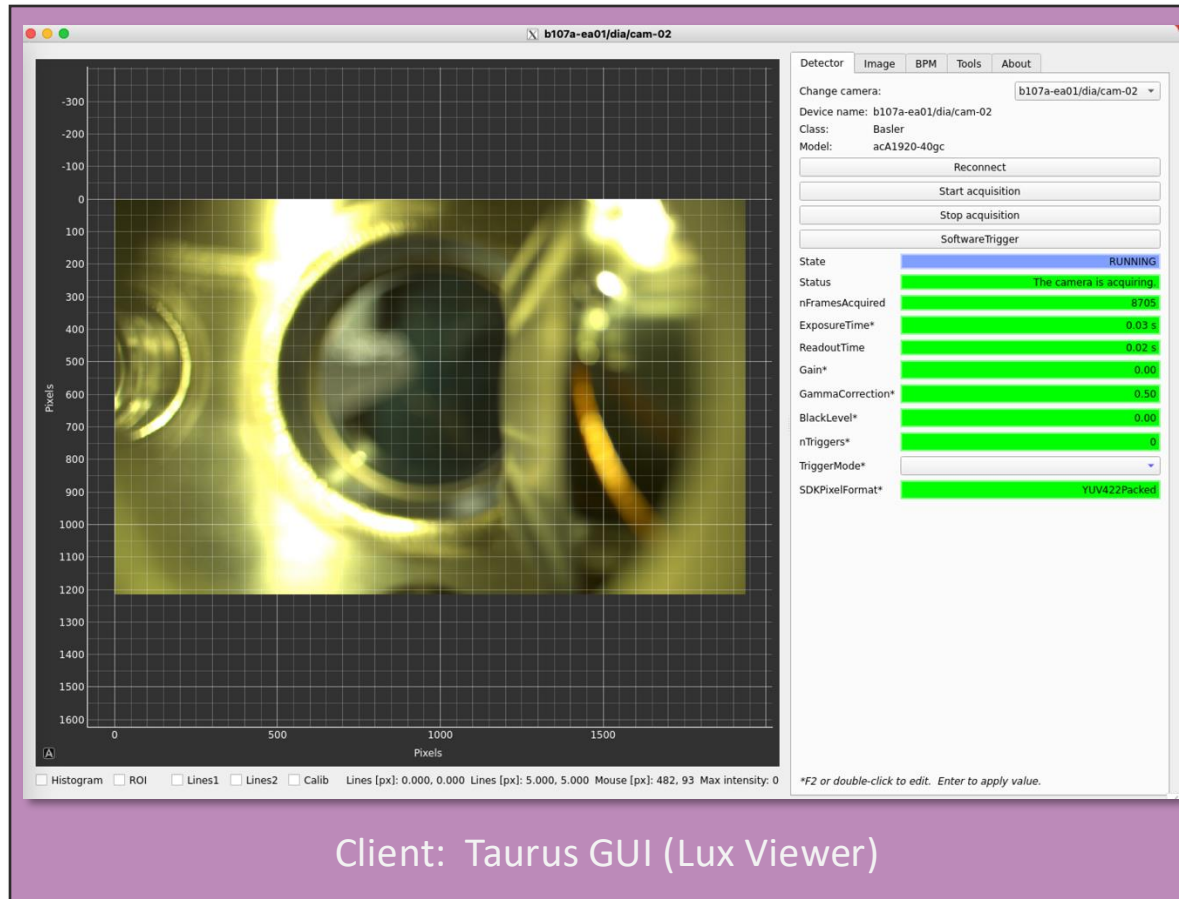
OpenTelemetry

Observability framework
Vendor agnostic
Open source
Many programming languages



Examples

GUI and Tango device for camera




The screenshot shows a software window titled "b107a-aa01/dia/cam-02". The main area displays a live camera feed of a mechanical part, overlaid on a coordinate grid. The vertical axis is labeled "Pixels" and ranges from -300 to 1600. The horizontal axis is also labeled "Pixels" and ranges from 0 to 1500. Below the feed, there are checkboxes for "Histogram", "ROI", "Lines1", "Lines2", and "Calib". A status bar at the bottom indicates "Lines [px]: 0.000, 0.000 Lines [px]: 5.000, 5.000 Mouse [px]: 482, 93 Max intensity: 0".

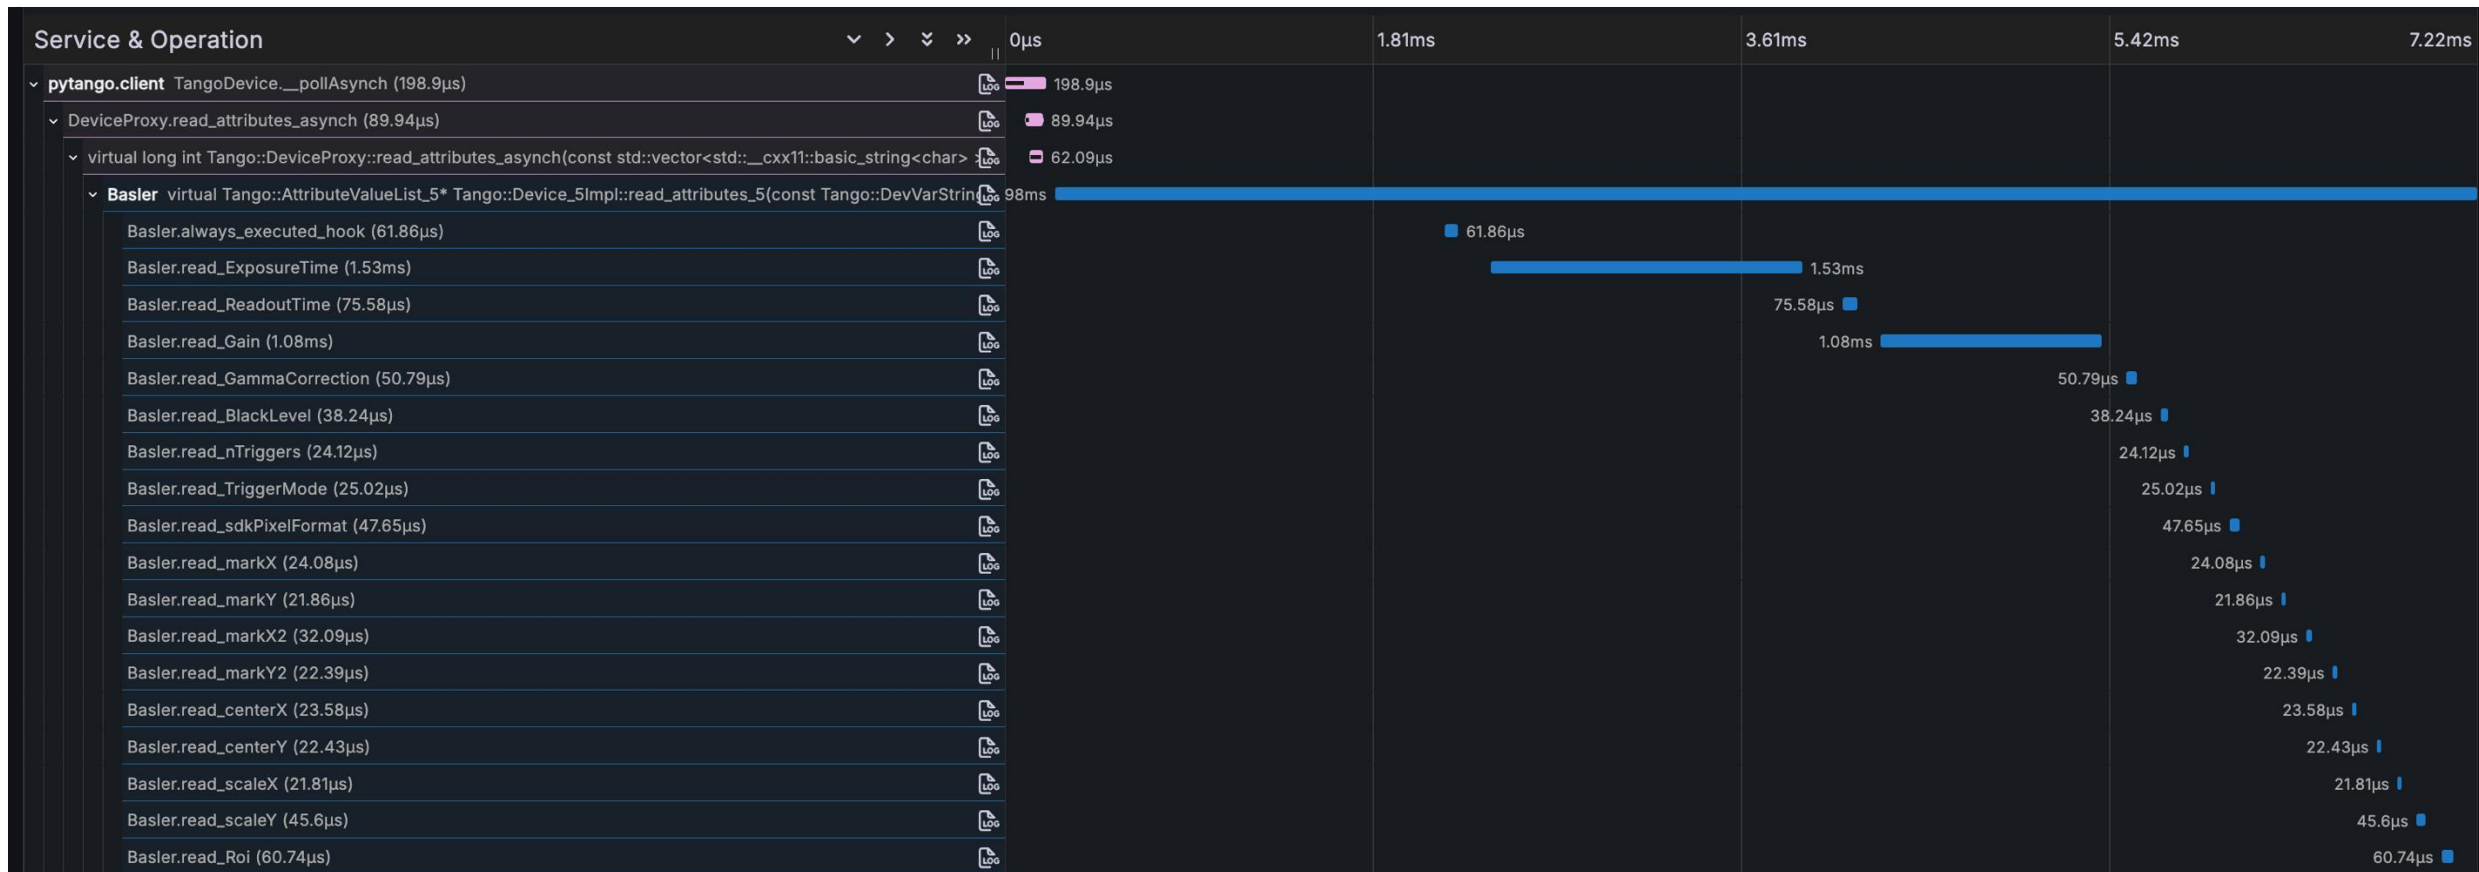
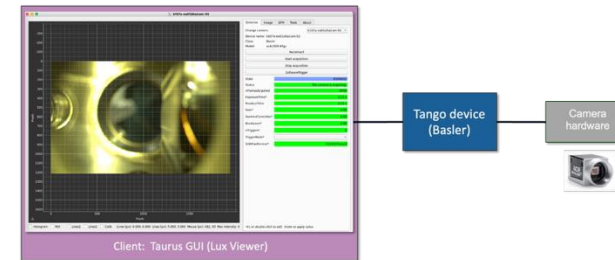
On the right side of the window is a control panel with the following elements:

- Detector: Image BPM Tools About
- Change camera: b107a-aa01/dia/cam-02
- Device name: b107a-aa01/dia/cam-02
- Class: Basler
- Model: acA1920-40gc
- Buttons: Reconnect, Start acquisition, Stop acquisition, SoftwareTrigger
- State: RUNNING
- Status: The camera is acquiring.
- nFramesAcquired: 8705
- ExposureTime*: 0.03 s
- ReadoutTime: 0.02 s
- Gain*: 0.00
- GammaCorrection*: 0.50
- BlackLevel*: 0.00
- nTriggers*: 0
- TriggerMode*: [dropdown menu]
- SDKPixelFormat*: YUV42Packed

Client: Taurus GUI (Lux Viewer)



Trace detail



User interface: [Grafana Tempo](#)

Span detail

The screenshot displays a distributed tracing interface with a dark theme. At the top, a 'Span Filters' section shows '22 spans' and navigation buttons for 'Prev' and 'Next'. A tree view on the left shows a hierarchy of spans: a root span for 'virtual long int Tango::DeviceProxy::read_attributes_async' (62.09µs), a child span for 'Basler' (98ms), and two sub-spans under 'Basler': 'Basler.always_executed_hook' (61.86µs) and 'Basler.read_ExposureTime' (1.53ms). The 'Basler.read_ExposureTime' span is selected and its details are shown in a large panel on the right. This panel includes the span name, service name ('Basler'), duration (1.53ms), start time (2.38ms), kind ('server'), and status ('unset'). It also lists the library name ('tango.python.server') and version ('10.0.2'). Below this, there is a 'Logs for this span' button and two sections: 'Span Attributes' and 'Resource Attributes'. The 'Span Attributes' section lists 'code.filepath', 'code.lineno', 'thread.id', and 'thread.name'. The 'Resource Attributes' section lists 'host.name', 'service.instance.id', 'service.name', 'service.namespace', 'telemetry.sdk.language', 'telemetry.sdk.name', and 'telemetry.sdk.version'. At the bottom right of the details panel, the 'SpanID' is shown as '51342392d3ce55f6'. The bottom of the interface shows a partial view of the next span, 'Basler.read_ReadoutTime' (75.58µs).

Span Filters 22 spans Prev Next

virtual long int Tango::DeviceProxy::read_attributes_async(const std::vector<std::__cxx11::basic_string<char>... 62.09µs

Basler virtual Tango::AttributeValueList_5* Tango::Device_5Impl::read_attributes_5(const Tango::DevVarStrin... 98ms

Basler.always_executed_hook (61.86µs)

Basler.read_ExposureTime (1.53ms)

Basler.read_ExposureTime Service: Basler Duration: 1.53ms Start Time: 2.38ms (11:34:16.584) Kind: server Status: unset Library Name: tango.python.server Library Version: 10.0.2

Logs for this span

Span Attributes

- code.filepath `"/opt/conda/envs/Basler/lib/python3.11/site-packages/Basler/Basler.py"`
- code.lineno `881`
- thread.id `"0x7fa6017fa700"`
- thread.name `"Dummy-6"`

Resource Attributes

- host.name `"bec-1"`
- service.instance.id `"b107a-ea01/dia/cam-02"`
- service.name `"Basler"`
- service.namespace `"tango"`
- telemetry.sdk.language `"python"`
- telemetry.sdk.name `"opentelemetry"`
- telemetry.sdk.version `"1.29.0"`

SpanID: 51342392d3ce55f6

Basler.read_ReadoutTime (75.58µs) 75.58µs

Filtering

▼ A (tempo-okd-test) [?] [📄] [👁] [🗑] [⋮]

Query type: Search | TraceQL | Service Graph Import trace

Service Name ⓘ = ▾ Basler × ▾ × ▾

Span Name ⓘ = ▾ Select value ▾

Status = ▾ Select value ▾

Duration ⓘ span ▾ > ▾ 10ms < ▾ e.g. 100ms, 1.2s

Tags span ▾ Select tag ▾ = ▾ Select value ▾

`{resource.service.name="Basler" && duration>10ms}` Edit in TraceQL

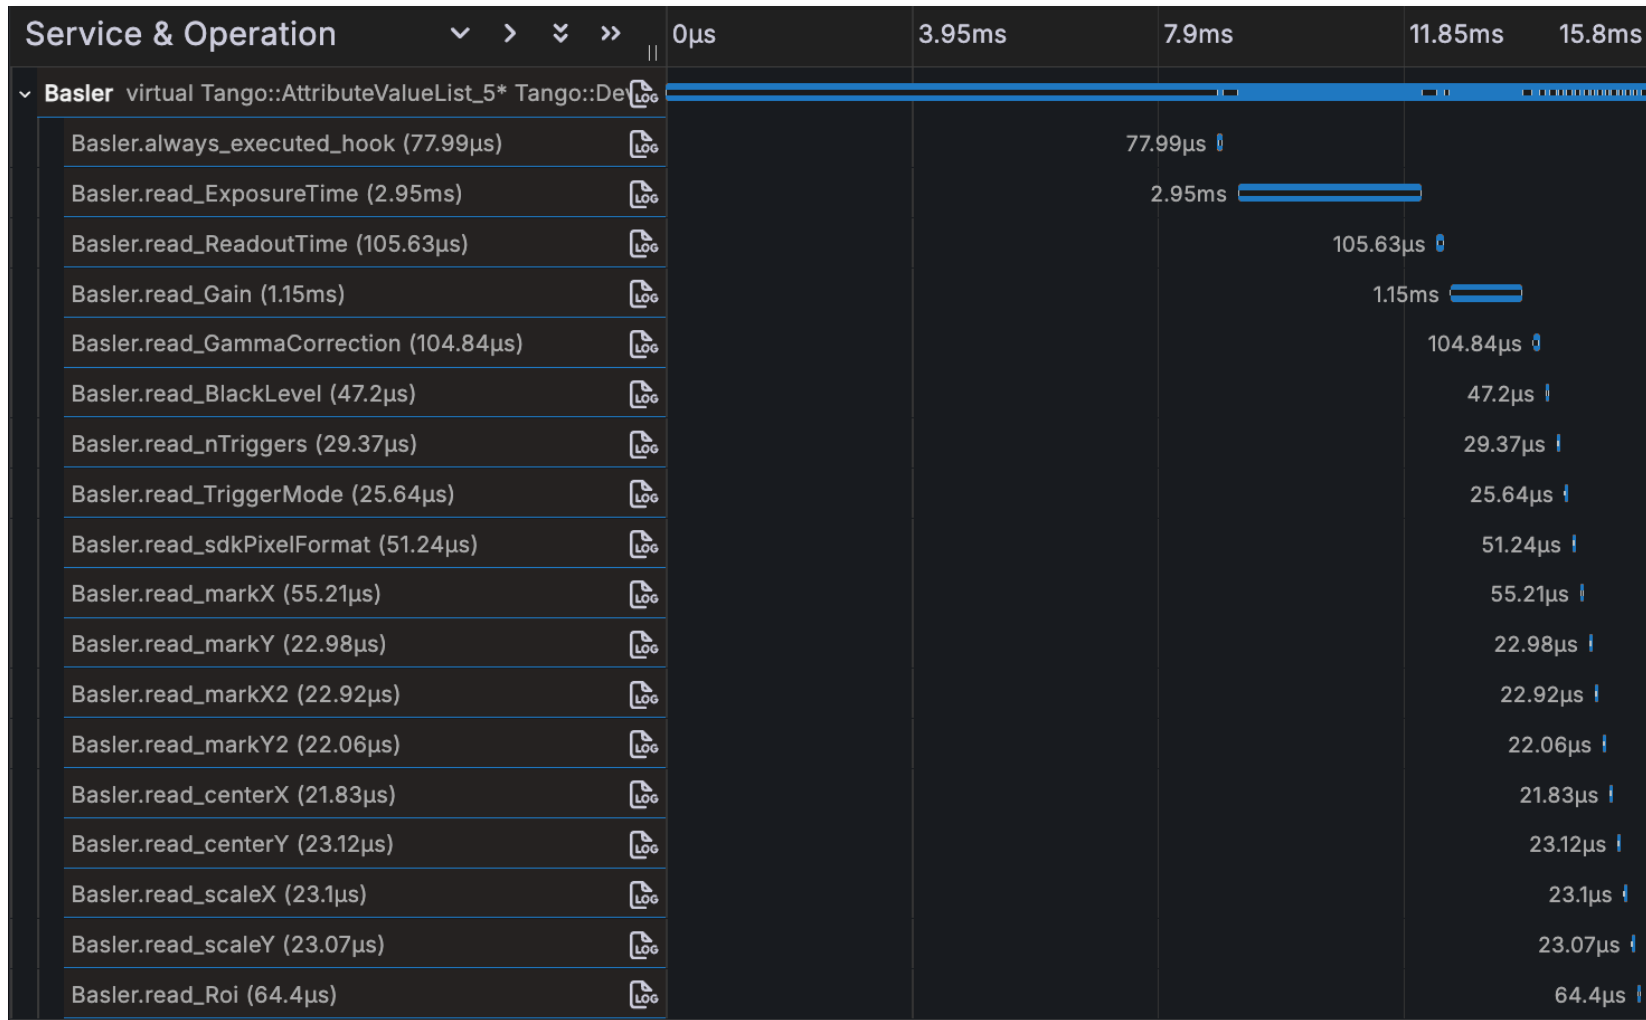
> Options Limit: 20 Spans Limit: 3 Table Format: Traces

+ Add query | 🔄 Query history | ⓘ Query inspector

Table - Traces

	Trace ID	Start time	Service	Name	Duration
>	516ae0ad947812d0520...	2025-03-31 11:47:32	Basler	virtual Tango::Attribute	10 ms
>	3b92629b059e004e4e...	2025-03-31 11:47:23	Basler	virtual Tango::Attribute	15 ms

Slow trace



Logs

The image displays two side-by-side Grafana dashboards. The left dashboard, titled 'tempo-rancher-obs-test', shows a trace for 'pytango.client: TaurusCommandButton._onClicked' with a duration of 15.56ms. The trace is visualized as a horizontal bar chart with 7 spans. The right dashboard, titled 'loki-rancher-obs-test', shows the logs for this trace. The query used is '{service_name="Basler", service_namespace="tango"}'. The logs are displayed in a table format, showing two log entries:

```
> 2025-04-07 12:00:06.318 Arming the camera.
> 2025-04-07 12:00:06.318 Destination Filename is empty, no data will be saved
```

An orange box highlights these two log entries, and an orange arrow points from the 'Basler' span in the trace to the first log entry. The trace spans are as follows:

Service & Operation	Start	End
pytango.client TaurusCommandButton._onClicked	0µs	15.56ms
Connection.command_inout (15.45ms)	0µs	15.45ms
virtual Tango::DeviceData Tango::Connection	0µs	15.45ms
Basler virtual CORBA::Any* Tango::Device	0µs	15.45ms
Basler.always_executed_hook (57.34µs)	11.67ms	11.72ms
Basler.is_Arm_allowed (25.74µs)	11.67ms	11.72ms
Basler.Arm (14.43ms)	11.67ms	13.11ms















Process info

```
$ export OTEL_EXPERIMENTAL_RESOURCE_DETECTORS=process  
$ python
```

Logs for this span

> Span Attributes: code.filepath = <stdin> | code.lineno = 2 | thread.id = 0x7fa71538c140 | thread.name = MainThread

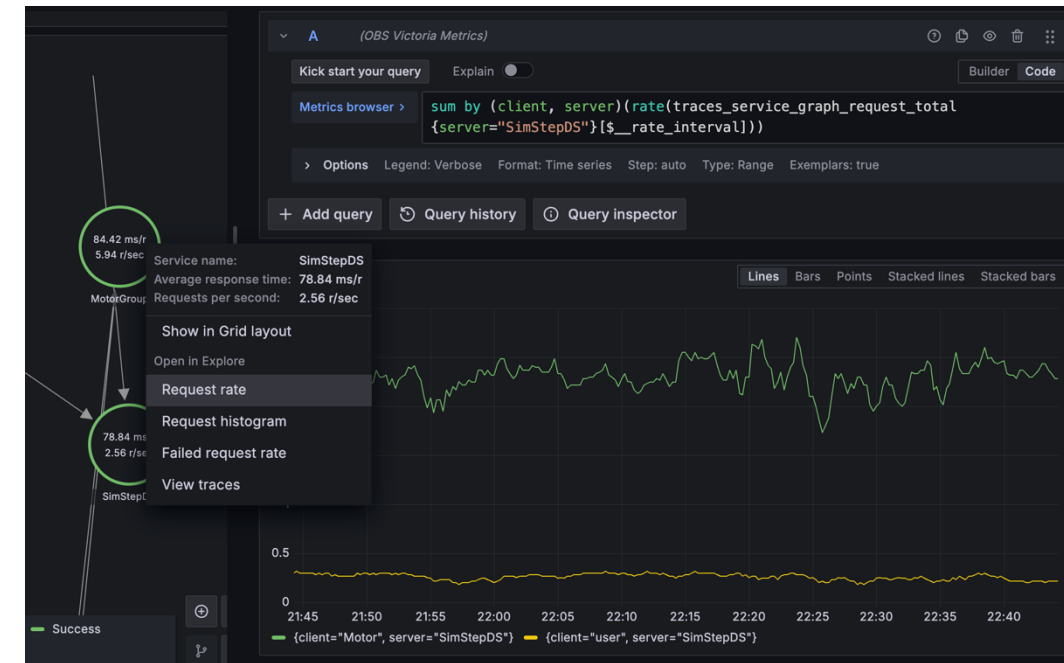
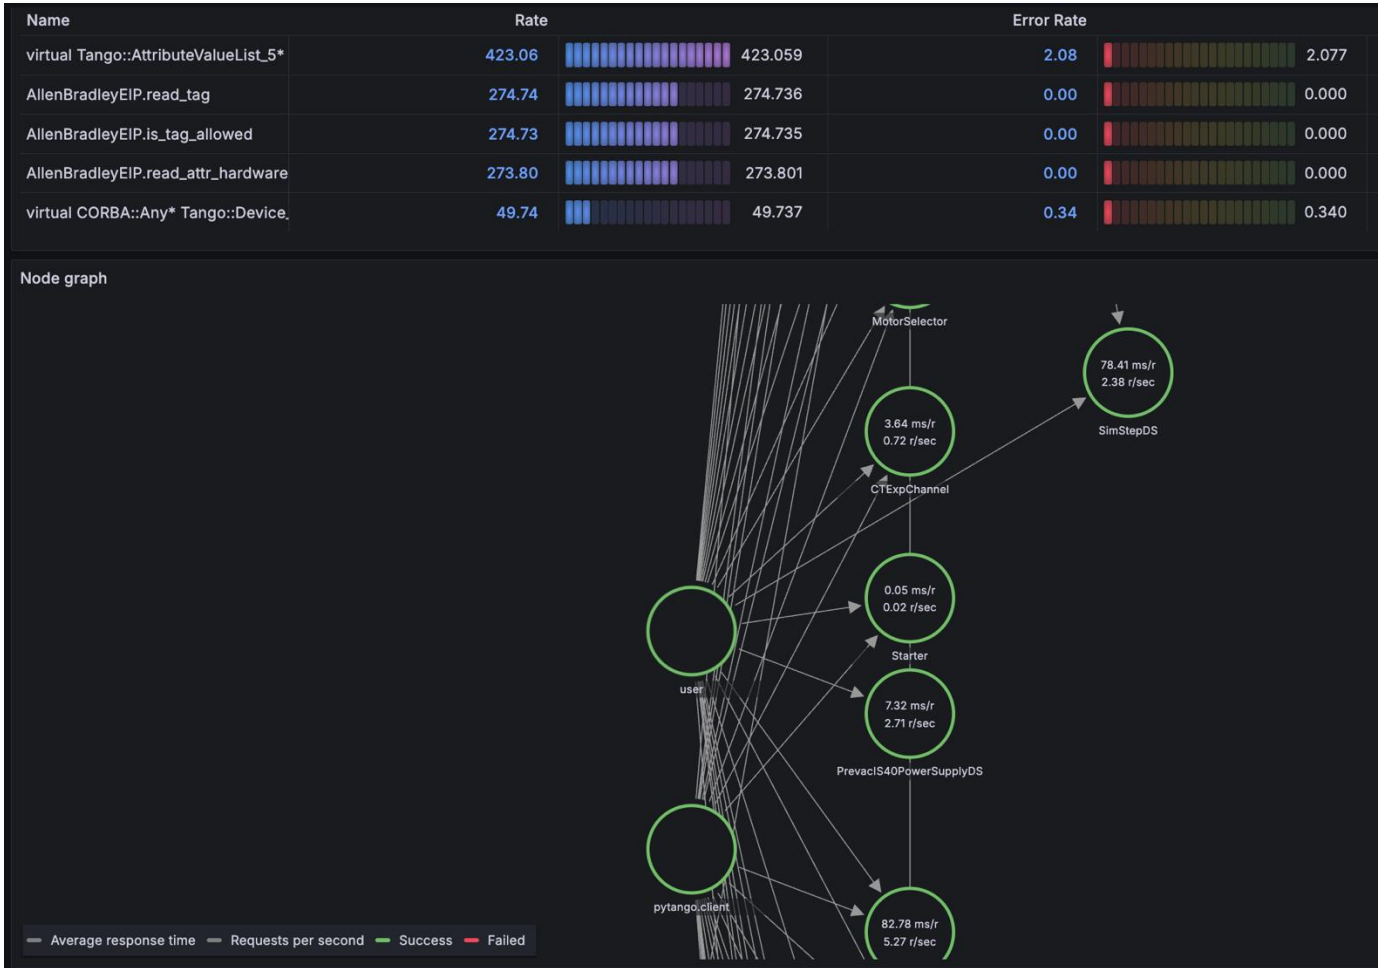
▼ Resource Attributes

host.name	"bc-0"	
process.command	""	
process.command_args	[""]	
process.command_line	""	
process.executable.name	"/opt/conda/envs/sardana/bin/python"	
process.executable.path	"/opt/conda/envs/sardana/bin"	
process.owner	"tango"	
process.parent_pid	331770	
process.pid	333402	
process.runtime.description	"3.11.11 packaged by conda-forge (main, Dec 5 2024, 14:17:24) [GCC 13.3.0]"	
process.runtime.name	"cpython"	
process.runtime.version	"3.11.11"	
service.name	"pytango.client"	
service.namespace	"tango"	

Time series from traces

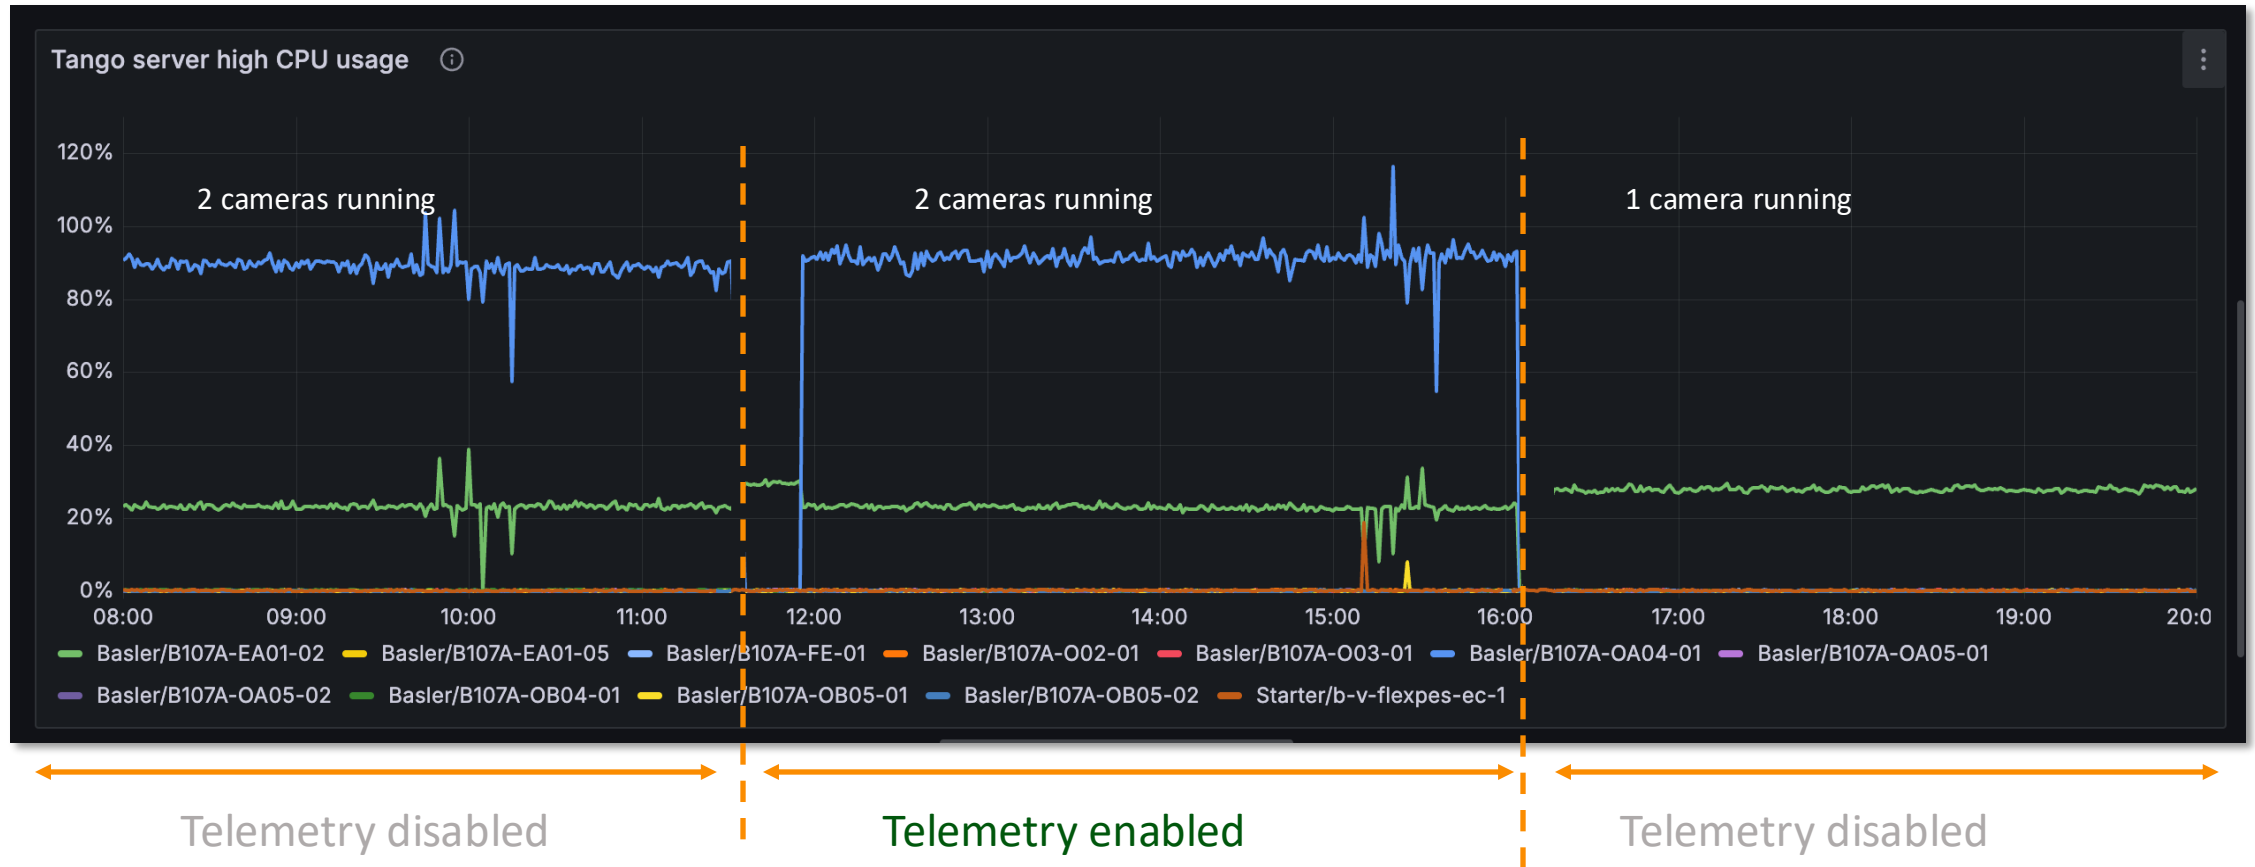


Service graph



Performance and resources

CPU usage (application)



Backend

Tested with 1 beamline, 900 Tango devices

Resources (2 replicas each for *Grafana Tempo* and *Loki*)

8 CPU cores

16 GB RAM

Storage

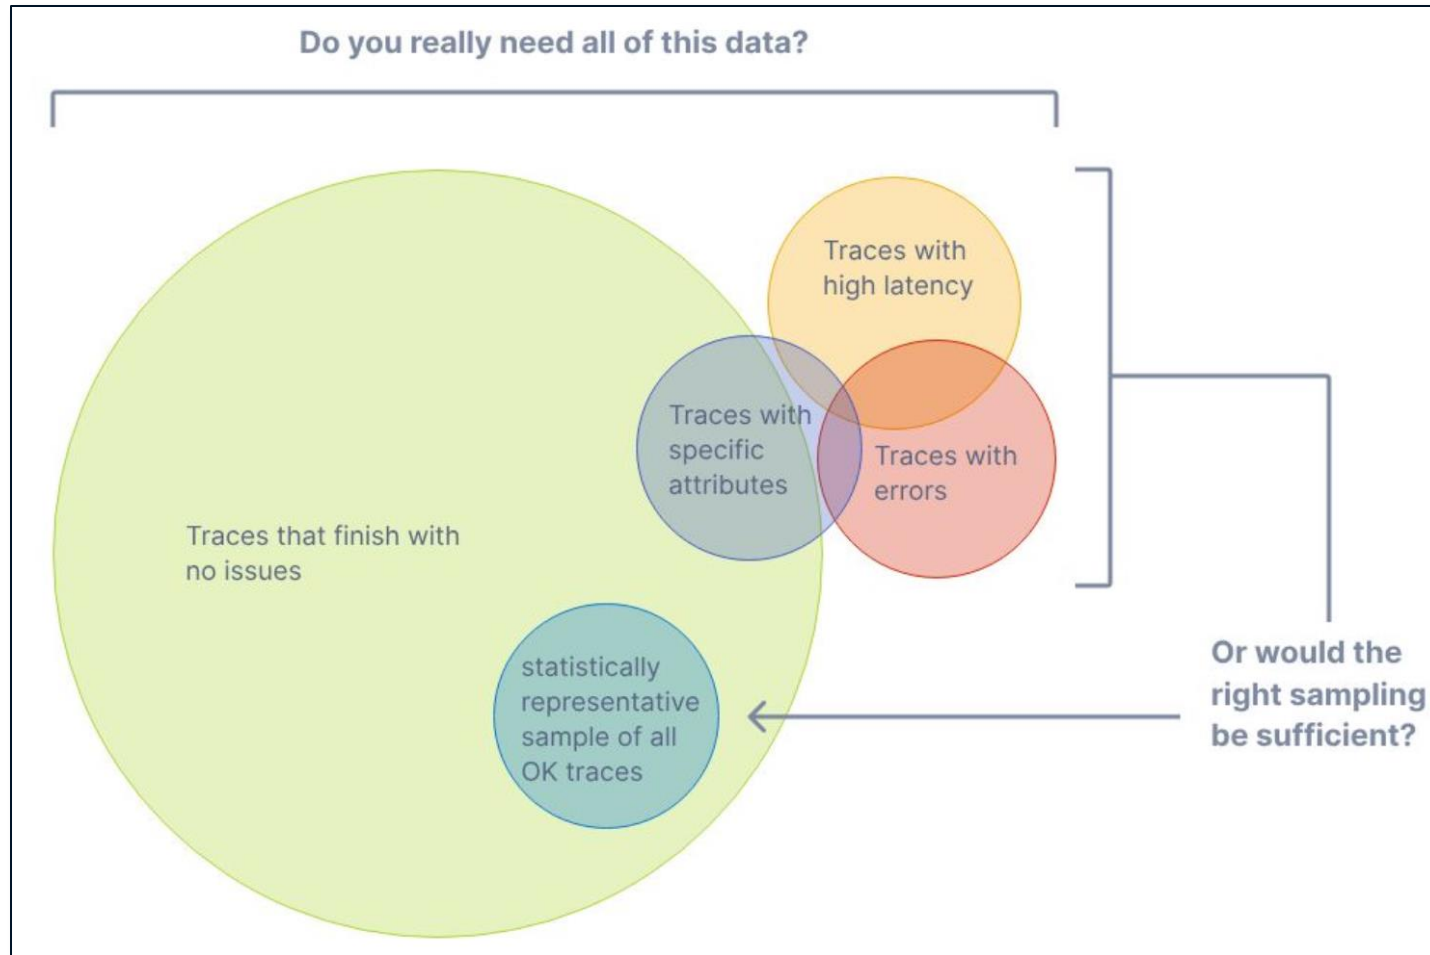
50 GB/day (only keeping 150 GB)

S3 storage

Network

3k spans/second => ~230 HTTP requests/sec

Sampling



Source: <https://opentelemetry.io/docs/concepts/sampling/>

Conclusion

Conclusion

Very powerful feature in Tango 10

Negligible performance impact on apps

Backend needs “a lot” of compute resources

Trace sampling will be needed to reduce trace volume

Room for improvement

More details in the paper:

[WEAG006](#)

BETTER SOFTWARE OBSERVABILITY USING TANGO CONTROLS WITH OPENTELEMETRY - EXPERIENCE AT MAX IV
A. F. Joubert*, L. Zhu, B. Bertrand, MAX IV, Lund, Sweden

Abstract
Distributed software systems are complex and the interactions across multiple machines can be difficult to debug and monitor. Log messages are not enough for observability. We need more information about the communication between applications, how each one is executing, and its internal state. In practice, applications can be made more observable using software frameworks such as OpenTelemetry. The Tango Controls framework has built-in support for OpenTelemetry in C++ and Python since version 10.0.0. We are using it operationally at the MAX IV synchrotron. We provide examples of the traces, trends, and other data available when running at scale on a beamline with hundreds of devices. We report on the compute and performance impact for client and server software applications, as well as practical issues. For the backend servers that ingest and query the telemetry data (running Grafana Tempo for traces and Grafana Loki for logs) we report on the compute resources required.




Figure 1: Example of a distributed trace across three applications (client script and two servers). The client script commands the Leader server to turn on a FoLLower server. The vertical axis shows a simplified call stack, and the horizontal axis shows the elapsed time since the first call.

OpenTelemetry
OpenTelemetry [6] is an open-source and vendor-agnostic software framework that provides a practical way to realise

Thank you!

Anton Joubert

anton.joubert@maxiv.lu.se

